

PROFESSIONAL RESPONSIBILITY IN THE AGE OF AI: DELEGATING THE WORK, NOT THE RESPONSIBILITY

Haris Hamidović^{ID}

MKF/MKD EKI Sarajevo, Sarajevo, Bosnia and Herzegovina

<https://orcid.org/0000-0002-1296-5008>

ARTICLE INFO



● Open access

JEL Category:

M15, O33

Keywords:

artificial intelligence

professional

responsibility

professional judgment

engineering ethics

human oversight

AI governance

ABSTRACT

Artificial intelligence (AI) is increasingly becoming part of professional practice. Engineers, auditors, cybersecurity professionals, expert witnesses, and other knowledge professionals can use AI to generate text, analyse information, produce software, identify patterns, develop alternatives, and support professional decisions. While these capabilities can increase professional capacity and efficiency, they also raise a fundamental question: when professionals delegate work to an AI system, can they also delegate professional responsibility? This paper argues that it cannot. AI may perform or assist with professional tasks, but its use does not, by itself, transfer professional responsibility for the resulting decision, opinion, design, recommendation, or action. This principle is not entirely new. Earlier professional debates concerning computer programs and computer-aided design addressed similar questions regarding technological assistance, professional competence, direction and control, and responsibility. Contemporary engineering ethics and AI governance frameworks reinforce the importance of competence, human oversight, accountability, risk management, and critical evaluation. Building on these principles, this paper proposes two complementary concepts: Delegable Work vs. Non-Delegable Professional Responsibility and an AI-Assisted Professional Judgment Model consisting of five stages: Define, Assist, Evaluate, Decide, and Account. The paper further proposes that professional AI literacy should be understood not merely as the ability to operate AI tools, but as the ability to understand, evaluate, challenge, and appropriately govern their outputs. The paper argues that increasing technological capability should not reduce professional responsibility. Instead, it should increase the importance of professional judgment, competence, critical reasoning, meaningful human oversight, and accountability.

Address of the author:

Haris Hamidović

[✉ haris.hamidovic@eki.ba](mailto:haris.hamidovic@eki.ba)

Received: 29.09.2026

Revised: 14.10.2026

Accepted: 16.10.2026

Available online: 16.10.2026



1 INTRODUCTION

Artificial intelligence is rapidly changing the way professional work is performed. Engineers use AI-assisted design and analysis tools. Auditors increasingly encounter automated and AI-generated information. Cybersecurity professionals use AI to analyse threats and identify patterns. Expert witnesses and consultants may use AI for research, document analysis, and drafting. Educators increasingly use AI as both a teaching tool and an object of professional and ethical discussion.

These developments create substantial opportunities. At the same time, they raise an important professional question: When a professional delegates part of the work to an AI system, what happens to professional responsibility for the result?

The question may appear to be new because the technology is new. However, the underlying professional issue is considerably older. Earlier generations of professionals faced similar questions when computer programs, automation, and computer-aided design became integrated into professional practice. Parker, Swope, and Baker's *Ethical Conflicts in Information and Computer Science, Technology and Business* included cases involving professional reliance on computer programs and questions concerning responsibility for results produced through technological systems. A central principle associated with those discussions was that authority may be delegated without transferring responsibility. (Parker et al., 1990)

The same question has now re-emerged with artificial intelligence. In *Looking Back to Move Forward*, I argued that while technological contexts change, many fundamental professional responsibilities remain remarkably stable, including responsibility toward society, professional competence, privacy, security, ethical conduct, and accountability. (Hamidovic, 2026a) In *Delegating the Work, Not the Responsibility*, I applied this principle directly to AI: professionals may delegate tasks to AI, but they cannot delegate their professional judgment and responsibility for the resulting work. (Hamidovic, 2026b)

This paper develops these two arguments into a broader position: AI may perform or assist with professional work, but the use of AI does not, by itself, transfer professional responsibility for the resulting professional decision, opinion, design, recommendation, or action. The objective is not to argue against AI adoption. Rather, the objective is to establish a professional framework within which AI can be used productively without confusing technological delegation with the delegation of professional accountability.

1.1 Original Contribution and Scope of the Paper

The contribution of this paper is not the introduction of the general principle that humans should remain responsible for the use of artificial intelligence. This principle is already reflected in professional codes of ethics and established AI governance frameworks. Rather, the paper develops a profession-centred conceptual and procedural approach to translating this principle into professional practice.

Recent work has examined the non-delegability of human judgment from a broader philosophical perspective. Torrecilla-Pinero (2026), for example, develops a subject-preserving framework for AI governance, arguing that algorithmic assistance may expand human deliberation but should not displace the human subject who exercises judgment. This perspective reinforces the distinction between using AI to support a decision and allowing its output to function as a substitute for human judgment. The present paper approaches this issue from a complementary, profession-centred perspective. Rather than primarily grounding the non-delegability of judgment in an account of personhood, it focuses on how professional responsibility can be exercised throughout an AI-assisted task. Its proposed contribution is the integration of the distinction between delegable work and non-delegable professional responsibility with a five-stage model intended to structure professional practice.

The contribution has two interrelated components. First, the paper explicitly distinguishes between the delegation of work that AI systems can perform or support and the non-delegable professional responsibility for evaluating, accepting, modifying,

or rejecting the resulting output. Second, it operationalizes this distinction through a five-stage AI-Assisted Professional Judgment Model: Define, Assist, Evaluate, Decide, and Account.

The conceptual distinction establishes what must remain with the professional, while the five-stage model describes how that responsibility can be exercised throughout an AI-assisted professional task. Their combination provides a practical structure for connecting professional judgment, meaningful human oversight, competence, and accountability. The proposed approach is intended to complement, rather than replace, existing professional ethics principles and AI governance frameworks.

2 METHODOLOGICAL APPROACH AND SCOPE

This paper adopts a conceptual analysis and normative argumentation approach to examine professional responsibility in AI-assisted professional practice. Rather than conducting an empirical investigation or a systematic literature review, it develops a profession-centred conceptual framework by critically examining established professional ethics principles, historical discussions of technological delegation, contemporary AI governance frameworks, and relevant academic literature.

The sources considered include professional codes of ethics, engineering ethics cases, academic publications on professional responsibility and automation, and normative and governance documents addressing AI risk management, accountability, and human oversight. Sources are considered relevant to the extent that they contribute to one or more of the following analytical dimensions: (1) the distinction between task execution and responsibility for professional outcomes; (2) professional competence and independent judgment; (3) human oversight and the limitations of automated decision support; and (4) the allocation of responsibilities among professionals, organizations, and AI system providers. Historical sources are used to identify established principles of professional practice, while contemporary sources are examined to determine how these principles relate to the capabilities and risks of AI systems.

The analysis follows a comparative and interpretive logic. Historical discussions of computer programs and computer-aided design are examined in relation to contemporary AI governance principles, with particular attention to the continuity of professional duties despite changes in the technologies used to perform professional work. The comparison distinguishes enduring normative principles, such as competence, direction and control, professional judgment, and accountability, from the practical challenges associated with AI-generated outputs, including uncertainty, automation bias, and the need for meaningful human oversight.

The proposed AI-Assisted Professional Judgment Model was developed by synthesizing these analytical dimensions into a sequence of professional activities and responsibilities. The distinction between delegable work and non-delegable professional responsibility provides the model's conceptual foundation. The five stages—Define, Assist, Evaluate, Decide, and Account—organize the process from defining the professional task and determining the appropriate role of AI to critically evaluating its outputs, exercising independent judgment, and accounting for the resulting professional outcome. The stages represent a conceptual and normative synthesis of the principles discussed in the paper rather than the results of empirical testing or a validated procedural standard.

This approach has limitations. The paper does not report a systematic review conducted under a predefined search protocol, nor does it empirically assess how professionals behave when using AI or whether the proposed model improves decision quality, reduces automation bias, or strengthens accountability. The analysis is therefore intended to provide a conceptual contribution and a basis for further investigation, rather than evidence of the model's effectiveness in practice. Future research could evaluate the model through expert interviews, case studies, scenario-based experiments, or field studies across different professional domains, examining its usability, applicability, and potential contribution to meaningful human oversight and professional accountability.

3 TECHNOLOGY CHANGES; PROFESSIONALISM ENDURES

Professional practice has always evolved with technology. Manual drafting gave way to computer-aided design. Manual calculations were supplemented by spreadsheets and specialized software. Auditors adopted increasingly sophisticated information systems and data analytics. Cybersecurity professionals moved from largely manual monitoring toward increasingly automated detection and response.

AI represents another significant step in this progression.

In my comparison of Robert Ayres's discussion of professional issues in computing with contemporary AI-related competencies, I found a clear continuity in the underlying professional concerns (Ayres, 1999; Hamidovic, 2026a). Although the technologies and terminology have changed, questions concerning the public interest, responsible use of technology, privacy, security, professional competence, and accountability remain relevant.

The technologies are different.

The professional questions are not. This suggests that AI should not be approached exclusively as a technological issue. It is also an issue of professional ethics and professional practice.

The relevant questions therefore extend beyond:

- What can AI do?
- How much work can AI automate?
- How quickly can AI produce an output?

They also include:

- What should the professional delegate?
- What should the professional verify?
- What level of review is appropriate?
- When should an AI output be rejected?
- Can the professional explain the basis for relying on the output?
- Who remains accountable for the resulting decision?

The distinction between technological capability and professional responsibility is therefore fundamental.

4 FROM COMPUTER PROGRAMS TO ARTIFICIAL INTELLIGENCE

One of the strongest arguments for treating AI as part of an established professional-responsibility debate comes from earlier computing technologies.

Parker et al. documented ethical conflicts arising when professionals relied on computer programs developed by others. (Parker et al., 1990) The underlying concern was already recognizable: if a professional's work depends upon a technological system that the professional did not create or fully control, does responsibility for the outcome shift to the technology developer?

The historical answer provides a useful foundation for AI.

A similar question arose in engineering with computer-aided design.

In NSPE Board of Ethical Review Case 90-6, the Society considered whether engineers could sign and seal documents produced using a CADD system. The case distinguishes between the technological production of the document and the engineer's professional responsibility for the work. NSPE states that the case established the ethical acceptability of signing and sealing CADD-produced documents where the work is under appropriate professional direction and control. (National Society of Professional Engineers [NSPE], 1990)

The significance of this precedent is not limited to CADD.

It demonstrates that professional responsibility does not require the professional to personally execute every technical operation.

Professionals have always relied on:

- tools;
- assistants;
- software;
- specialized experts;
- automation;
- calculations;
- models; and
- technological systems.

What remains essential is appropriate professional competence, direction, control, evaluation, and responsibility for the resulting work.

The transition from CADD to AI therefore represents an important technological development, but not necessarily a fundamental change in the underlying principle of professional accountability.

The relevance of this principle to contemporary AI-assisted engineering practice is further illustrated by NSPE Board of Ethical Review Case 24-2, which specifically addresses the use of artificial intelligence in engineering practice. The case provides a profession-specific perspective on the ethical considerations that arise when engineers use AI tools in performing their professional duties. It complements the earlier CADD case by placing the relationship between technological assistance and professional responsibility in the context of contemporary AI systems. (National Society of Professional Engineers [NSPE], 2024)

4.1 Contemporary AI Governance and Professional Responsibility

Contemporary AI governance frameworks provide an important context for understanding professional responsibility in AI-assisted work. However, they address this issue from different, complementary perspectives. The EU Artificial Intelligence Act establishes legal requirements for specific categories of AI systems, including risk management, transparency, and human oversight for high-risk systems. Article 14 is particularly relevant because it addresses effective human oversight, the ability to understand system capabilities and limitations, and the risk of automation bias (European Parliament and Council of the European Union, 2024).

The NIST AI Risk Management Framework approaches trustworthy AI through four functions: Govern, Map, Measure, and Manage. It emphasizes organizational governance, contextual risk identification, risk assessment, and ongoing risk management throughout the AI system lifecycle (NIST, 2023). Similarly, the OECD AI Principles emphasize human-centred values, transparency, robustness, safety, and accountability, while recognizing that responsibility should reflect the roles and contexts of different AI actors (OECD, 2024). ISO/IEC 42001:2023 complements these approaches by specifying requirements for establishing and

continually improving an organizational AI management system (ISO, 2023).

These frameworks demonstrate that responsible AI cannot be reduced to the behaviour of an individual user. It also depends on organizational arrangements, technical safeguards, documented processes, and the responsibilities of providers and deployers. Nevertheless, the existence of governance arrangements does not, by itself, establish whether a professional has exercised adequate judgment when relying on a particular AI-generated result.

This distinction is central to the present paper. Organizational frameworks address how AI-related risks and responsibilities should be governed, whereas the proposed approach focuses on the professional's judgment in relation to a specific AI-assisted task and its outcome. It does not assume that all responsibility belongs exclusively to the individual professional. Rather, it distinguishes the professional's responsibility for a particular decision from the wider responsibilities of organizations, AI providers, and other actors. The five-stage model subsequently developed in this paper provides a structured way to connect the use of AI with evaluation, decision-making, and professional accountability.

5 THE PRINCIPLE OF NON-DELEGABLE PROFESSIONAL JUDGMENT

Professional judgment represents the critical boundary between AI assistance and professional responsibility. An AI-generated result may be plausible without being correct. It may be technically sophisticated without being appropriate. It may be factually accurate in isolation while being inappropriate in context. It may be useful for one purpose while creating unacceptable risks for another.

A professional using AI should therefore consider, as appropriate:

1. Is the output accurate?
2. Is the underlying information appropriate and sufficiently complete?
3. Are the assumptions reasonable?
4. Is the method suitable for the intended purpose?

5. Is the result consistent with applicable professional standards?
6. What are the limitations of the AI system?
7. What are the consequences if the result is wrong?
8. What additional verification is required?
9. Are privacy, cybersecurity, legal, regulatory, or ethical requirements affected?
10. Does the professional have sufficient competence to evaluate the output?

This distinction has direct implications for AI literacy.

AI literacy should not be reduced to knowing how to operate an AI system or formulate effective prompts.

For professionals, AI literacy should include the ability to:

- understand capabilities and limitations;
- identify potential failure modes;
- challenge plausible outputs;
- verify important results;
- recognize contextual limitations;
- understand relevant privacy and security implications;
- assess risks; and
- determine when AI should not be relied upon.

As I argued in *Delegating the Work, Not the Responsibility*, a professional cannot simply accept an AI-generated result because it appears plausible. Professional skepticism, independent assessment, and understanding of AI limitations remain necessary.

5.1 Automation Bias and the Limits of Human Verification

An important challenge to meaningful professional judgment is automation bias, understood as the tendency to rely excessively on automated recommendations or to accept system outputs without sufficient independent evaluation. Research on decision-support systems has shown that automation can introduce new forms of error when users fail to detect incorrect recommendations or become less vigilant in monitoring system outputs (Goddard et al., 2012). More recent research on human–AI collaboration

similarly identifies automation bias as a continuing concern, influenced by factors such as trust in AI, task conditions, user expertise, and the demands of verification (Romeo & Conti, 2025).

Automation bias is particularly relevant to professional practice because AI-generated outputs may be coherent, detailed, and persuasive even when their assumptions, evidence, or conclusions are incorrect. Consequently, the mere presence of a professional reviewer does not establish that meaningful oversight has occurred. A review that consists primarily of confirming an apparently plausible result may preserve the appearance of human involvement while failing to provide effective independent judgment.

The distinction between delegable work and non-delegable professional responsibility therefore requires more than a formal approval step. It requires the professional to assess the suitability of the output, identify material uncertainties, seek independent verification where appropriate, and remain able to modify or reject the recommendation. The degree of verification should be proportionate to the potential consequences of error, the reliability of the available evidence, and the professional context.

Within the proposed five-stage model, the Evaluate and Decide stages are particularly relevant to mitigating automation bias. Evaluate requires critical examination rather than passive acceptance, while Decide preserves the professional's authority to accept, modify, supplement, independently verify, or reject the AI-generated output. The Account stage further requires the professional to explain the basis for relying on the result. These stages do not eliminate automation bias, but they provide a procedural structure for addressing its implications for professional judgment and accountability.

6 DELEGABLE WORK VS. NON-DELEGABLE PROFESSIONAL RESPONSIBILITY

The distinction between work and responsibility can be expressed through a two-layer conceptual model.

Table 1. *Delegable Work vs. Non-Delegable Professional Responsibility*

AI-assisted or delegable work	Non-delegable professional responsibility
Calculations	Determining whether the calculation is appropriate
Data analysis	Assessing whether data and method are suitable
Drafting and summarization	Verifying accuracy and contextual relevance
Code generation	Determining whether code is secure, reliable, and fit for purpose
Pattern identification	Determining the significance of the identified pattern
Preliminary design alternatives	Selecting and approving an appropriate design
Recommendations	Deciding whether to accept, modify, or reject the recommendation
Information retrieval	Assessing reliability and relevance
Scenario generation	Assessing assumptions, risks, and consequences

The first column does not mean that all listed activities should be delegated to AI.

Rather, it identifies activities that technology may potentially perform or support.

The second column identifies professional decisions that cannot be transferred merely because the first-column activity has been automated.

The distinction therefore is not:

- AI work vs. human work, but:
- delegable execution vs. non-delegable professional accountability.

This distinction avoids two opposite misconceptions.

The first is that professionals must personally perform every task to remain responsible. This would be inconsistent with established professional practice.

The second is that responsibility automatically follows a task once it is delegated to an AI system.

The appropriate relationship is: AI performs or assists → Professional evaluates → Professional decides → Professional remains accountable.

7 AI-ASSISTED PROFESSIONAL JUDGMENT MODEL

Based on the preceding distinction, this paper proposes an AI-Assisted Professional Judgment Model.

The model consists of five stages:

A. DEFINE

The professional defines:

- the problem;
- purpose;
- context;
- applicable requirements;
- constraints; and
- acceptable risk.

B. ASSIST

AI is used to perform or support appropriate tasks such as:

- analysis;
- calculation;
- drafting;
- research;
- classification;
- recommendations; or
- generation of alternatives.

C. EVALUATE

The professional examines the AI output in relation to:

- accuracy;
- relevance;
- assumptions;
- limitations;
- consistency;
- potential bias;
- security implications; and
- consequences.

D. DECIDE

The professional determines whether the output should be:

- accepted;
- modified;
- supplemented;
- independently verified; or
- rejected.

E. ACCOUNT

The professional remains able to:

- explain the decision;
- substantiate its basis;
- document relevant considerations; and
- accept responsibility for the resulting professional outcome.

The five stages are not intended to represent an entirely new set of ethical obligations. Rather, they organize established professional duties into a coherent sequence for AI-assisted work. The model connects the definition of the professional task with the use of AI, critical evaluation of its output, the exercise of independent professional judgment, and the ability to account for the resulting decision.

The relationship between the two contributions is fundamental to the proposed approach. The distinction between delegable work and non-delegable professional responsibility establishes the conceptual principle; the five-stage model translates that principle into a process that can guide professional practice. The model does not eliminate the need for organizational AI governance, domain-specific standards, or risk-based controls. Instead, it provides a professional judgment structure that can operate alongside them.

7.1 Practical Application of the AI-Assisted Professional Judgment Model

To illustrate how the proposed model can be applied in practice, consider a hypothetical scenario in which an internal auditor uses an AI system to identify potentially anomalous transactions in a financial dataset. Although the system may process large volumes of data and identify patterns that would be difficult to detect manually, its outputs do not, by themselves, establish that a transaction is erroneous, fraudulent, or evidence of misconduct.

Define. The auditor first defines the audit objective, the scope of the analysis, the relevant

criteria, and the consequences of an incorrect conclusion. The auditor also considers whether the available data are sufficiently complete, accurate, and appropriate for the intended purpose. This stage establishes the professional question that the AI system is expected to support, rather than allowing the system's capabilities to determine the audit objective.

Assist. The AI system is used to identify unusual transaction patterns, inconsistencies, or items that warrant further examination. Its role is to support the auditor's analysis by prioritising potential areas of concern. The use of AI does not transfer responsibility for defining the audit procedures or interpreting the significance of the results to the system.

Evaluate. The auditor critically examines the identified anomalies, including the underlying data, the criteria applied, and the limitations of the AI system. Relevant transactions are checked against source records, supporting documentation, applicable controls, and other available evidence. The auditor considers alternative explanations, including legitimate business reasons for unusual patterns, rather than treating an AI-generated flag as proof of an irregularity. The depth of verification is proportionate to the potential consequences of an incorrect assessment.

Decide. Based on the verified evidence and independent professional judgment, the auditor determines whether further investigation is necessary and whether the findings support an audit conclusion. The decision may differ from the AI system's initial indication. Where the evidence is insufficient or contradictory, the auditor may seek additional information, expand the testing, or refrain from drawing a definitive conclusion.

Account. The auditor documents the purpose for which AI was used, the relevant inputs and outputs, the verification procedures performed, the evidence considered, the rationale for the final decision, and any material limitations. This documentation supports the reviewability of the work and demonstrates how the professional conclusion was reached. The auditor remains accountable for the audit judgment and the conclusions communicated within the scope of the auditor's professional role.

For practical purposes, the documentation should be sufficient to explain how the AI-assisted result contributed to the professional conclusion and to enable appropriate review of the process. Depending on the nature of the task, the potential consequences of error, and applicable requirements, relevant records may include the purpose and context of AI use; the AI system used and its version, where available; material inputs and limitations affecting the analysis; verification procedures performed and evidence considered; identified uncertainties or deficiencies; the rationale for accepting, modifying, supplementing, independently verifying, or rejecting the output; and the final professional decision and the identity of the responsible professional, where appropriate. The extent of documentation should be proportionate to the risks and complexity of the task and should respect applicable confidentiality, privacy, security, and record-retention requirements. This approach is intended to support traceability and accountability without imposing an unnecessarily burdensome documentation process for every AI-assisted activity.

This example illustrates that the model does not prohibit the delegation of analytical work to AI. Instead, it distinguishes the use of AI to generate or prioritise information from the professional responsibility to evaluate evidence, make a justified decision, and account for the resulting conclusion.

Although the five stages are presented in a logical sequence, the model should not be interpreted as a strictly linear process. The professional may need to revisit an earlier stage when evaluation reveals incomplete data, unsuitable assumptions, insufficient evidence, or unreliable AI outputs. For example, deficiencies identified during the Evaluate stage may require the professional to return to Define and refine the task or obtain additional data before repeating the analysis. Similarly, if the evidence does not support a defensible conclusion, the professional may need to conduct further verification before proceeding to Decide. Iteration is therefore an integral feature of the model, allowing professional judgment to respond to uncertainty and newly identified risks.

8 THE POSITION OF THIS PAPER

The paper does not claim that AI should be "used responsibly." Such a proposition is already reflected in professional codes, AI governance frameworks, and international ethical guidance.

The more specific position advanced here is:

Professional responsibility should be treated as non-delegable even when professional work is increasingly delegated to AI systems.

This position has three principal propositions.

Proposition 1: Task delegation does not imply responsibility delegation

The ability of AI to perform professional tasks should not be interpreted as a transfer of professional accountability to the AI system, its developer, or its provider.

Proposition 2: Professional judgment is the critical human control point

The professional's essential role is not necessarily to perform every underlying task personally, but to determine whether the AI-assisted result is appropriate for the professional context and the consequences involved.

Proposition 3: Professional competence must evolve with technological capability

As AI becomes more capable, professional competence must increasingly include the ability to understand, evaluate, challenge, and appropriately govern AI-generated outputs.

The proposed approach builds on existing principles of responsible AI, including human oversight, accountability, transparency, risk management, and professional competence. In particular, UNESCO's Recommendation on the Ethics of Artificial Intelligence emphasizes human responsibility and accountability, while professional codes of ethics establish duties concerning competence, public welfare, and responsible professional conduct. These principles provide an essential normative foundation for the argument developed in this paper.

The added value of the proposed approach lies in its specific focus on the professional task and the decisions surrounding the use of AI-generated outputs. Rather than introducing a separate set of

general AI ethics principles, it connects the delegation of individual work activities with the professional's continuing obligation to evaluate the results and take responsibility for the resulting professional outcome.

The five-stage model provides a structured sequence through which these obligations can be considered in practice. Define establishes the purpose, context, constraints, and acceptable risk; Assist identifies the appropriate use of AI; Evaluate requires critical examination of the output; Decide preserves the professional's authority to accept, modify, verify, or reject it; and Account emphasizes the ability to explain and substantiate the decision and accept responsibility for the outcome.

Accordingly, the model is best understood as a complementary professional-practice framework rather than a substitute for existing AI governance frameworks. Organizational governance establishes policies, controls, roles, and risk-management arrangements, while the proposed model focuses on how an individual professional exercises judgment and accountability within those arrangements. Its application should therefore be proportionate to the professional context, the potential consequences of error, and applicable legal, regulatory, and organizational requirements.

These propositions lead to a broader principle.

The purpose of responsible AI in professional practice is not to keep humans performing tasks that machines can perform more efficiently. It is to ensure that technological delegation does not eliminate meaningful professional judgment and accountability.

The position can therefore be summarized in three short statements:

- Delegate tasks where appropriate.
- Retain judgment where necessary.
- Never delegate accountability merely because technology can perform the work.

8.1 Comparative Positioning and Added Value of the Proposed Model

The proposed AI-Assisted Professional Judgment Model should be interpreted in relation to existing approaches to AI governance, risk management,

and responsibility. These approaches share important objectives, but they differ in scope and primary purpose.

The EU AI Act establishes legal obligations for providers and deployers of AI systems within its scope, including specific requirements for human oversight of high-risk systems. The NIST AI Risk Management Framework provides a structured approach to identifying, assessing, and managing AI risks at the organizational and system levels. The OECD AI Principles articulate broader principles for trustworthy AI and emphasize accountability across the roles of different AI actors. ISO/IEC 42001 provides requirements for an organizational management system for AI. None of these instruments should be understood as a universal, step-by-step procedure for exercising professional judgment in every AI-assisted task.

The proposed model addresses this professional-practice perspective. Its purpose is not to reproduce existing governance requirements or introduce an alternative regulatory framework. Instead, it organizes the professional's responsibilities around five connected stages: defining the task and its constraints, using AI appropriately, critically evaluating the output, making an independent professional decision, and remaining accountable for the resulting outcome.

The distinction is therefore complementary rather than competitive. Organizational governance establishes the policies, roles, controls, and resources within which AI is used. The proposed model focuses on how a professional should exercise judgment within that environment. It also makes explicit the connection between delegating an activity and retaining responsibility for deciding whether its result can appropriately be relied upon.

This approach must nevertheless be understood within the limits of professional and organizational accountability. A professional's responsibility for a particular decision does not automatically remove the responsibilities of an organization for inadequate governance, a provider for relevant deficiencies in system design or documentation, or other actors for their own conduct. The applicable allocation of legal liability depends on the relevant law, contractual arrangements, the roles of the parties, and the circumstances of the case.

The allocation of responsibility in AI-assisted professional practice must therefore be understood as complementary rather than exclusive. The professional remains responsible for exercising appropriate judgment when evaluating and relying on AI-generated outputs, while the organization is responsible for establishing appropriate conditions, governance arrangements, and controls for AI use. AI system providers and other involved parties may also bear responsibilities arising from their respective roles, contractual obligations, and applicable legal requirements. Retaining professional responsibility for a particular decision does not eliminate or automatically transfer the responsibilities of other actors. The applicable allocation of legal liability depends on the relevant law, the roles and conduct of the parties, contractual arrangements, and the circumstances of the case.

The added value of the model therefore lies in the integration of a conceptual distinction and a practical decision sequence. The distinction establishes that technological delegation does not, by itself, transfer professional accountability; the five-stage model translates this principle into a structured approach to AI-assisted work. Whether the model improves professional practice, reduces automation bias, or produces better accountability outcomes remains a question for future empirical evaluation.

9 PROFESSIONAL COMPETENCE IN THE AGE OF AI

The emergence of AI does not make traditional professional competence less important.

It changes what professional competence must include.

The IEEE Code of Ethics, adopted in its revised form in June 2020, calls on IEEE members to uphold public safety and welfare, improve understanding of the capabilities and societal implications of emerging technologies including intelligent systems, maintain and improve technical competence, undertake technological tasks only when appropriately qualified or after disclosing relevant limitations, and seek and offer honest technical criticism and correct errors. (IEEE, 2020)

These principles are directly relevant to AI-enabled professional practice.

Professional AI competence can therefore be understood through at least five dimensions:

1. Domain competence. The professional must understand the subject matter to which AI is applied.
2. AI literacy. The professional must understand AI capabilities, limitations, uncertainty, and relevant failure modes.
3. Critical reasoning. The professional must be able to challenge plausible outputs rather than accepting them because they appear coherent.
4. Ethical and regulatory awareness. The professional must recognize implications involving privacy, security, intellectual property, discrimination, safety, accountability, and applicable legal requirements.
5. Governance awareness. The professional must understand how AI use fits within organizational policies, controls, risk management, and accountability arrangements.

This represents a shift from treating AI primarily as a productivity skill toward understanding AI as part of professional competence and responsibility.

10 HUMAN OVERSIGHT IS NOT A FORMALITY

Human oversight is sometimes interpreted simply as having a human somewhere in the process.

That is insufficient.

If a human merely accepts an AI-generated result without possessing the competence, information, authority, or time necessary to evaluate it, the existence of a human reviewer does not necessarily constitute meaningful oversight.

UNESCO's Recommendation on the Ethics of Artificial Intelligence identifies responsibility and accountability, transparency and explainability, and human oversight and determination among its core principles. It explicitly states that AI systems should not displace ultimate human responsibility and accountability. It also calls for auditability, traceability, impact assessment, audit, and due diligence mechanisms. (UNESCO, 2022)

Meaningful professional oversight should therefore include:

- appropriate competence;
- sufficient contextual information;
- authority to challenge or reject AI output;
- adequate time and resources for review;
- verification proportionate to risk;
- documentation where appropriate; and
- clear accountability for the resulting decision.

Human oversight should not become a ceremonial step in an automated process.

The relevant question is not merely: "Was a human involved?"

It is: "Was the human sufficiently competent, informed, empowered, and involved to exercise meaningful professional judgment?"

11 CONCLUSIONS

Artificial intelligence is changing professional practice at an unprecedented pace.

Yet the most important professional question it raises is not entirely new.

Earlier generations of professionals confronted computer programs, automation, and computer-aided design. They asked whether reliance on technological systems could change who remained responsible for professional work.

The principle remains relevant. Technology can change how work is performed without changing who remains professionally responsible for the decision to rely on the result. AI can draft, calculate, analyze, classify, generate, and recommend. But professional responsibility requires something more: the ability to determine whether output is appropriate, reliable, safe, ethical, and fit for purpose.

The challenge of the AI era is therefore not to preserve a world in which professionals perform every task themselves. Professional practice has always involved delegation, collaboration, tools, automation, and technological assistance.

The challenge is to preserve the distinction between delegating work and delegating responsibility.

The principal contribution of this paper is the integration of a conceptual distinction between delegable work and non-delegable professional responsibility with a five-stage model for exercising professional judgment in AI-assisted practice. While existing AI ethics and governance frameworks establish important principles of accountability and human oversight, the proposed approach organizes these principles around the professional task and the decisions required before an AI-generated result can be relied upon. Its value lies not in replacing established frameworks, but in providing a complementary structure for applying their underlying principles to professional judgment, decision-making, and accountability.

This paper proposes a simple professional principle:

Delegate tasks where appropriate. Retain judgment where necessary. Never delegate accountability merely because technology can perform the work.

Technology will continue to change. Professionalism must evolve with it. But professionalism should not disappear behind technology. Technology may change. The responsibility remains.

REFERENCES

- Ayres, R. (1999). *The essence of professional issues in computing*. Prentice Hall.
- European Parliament & Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, 19(1), 121–127. <https://doi.org/10.1136/amiajnl-2011-000089>

- Hamidovic, H. (2026a, September). Looking back to move forward – Teaching future engineers that professionalism outlasts technology. *IEEE SSIT Insights*. https://engage.ieee.org/ES-TA-1033-2026-09-SSIT-Newsletter_SSIT-Insights-Page-2026-09.html
- Hamidovic, H. (2026b, October). Delegating the work, not the responsibility: Professional judgment in the age of AI. *IEEE SSIT Insights*. <https://engage.ieee.org/SSIT-Newsletter-Insights-October-2026.html>
- IEEE. (2020, June). *IEEE Code of Ethics* (IEEE Policies, Section 7.8). <https://www.ieee.org/about/corporate/governance/p7-8.html>
- International Organization for Standardization. (2023). *ISO/IEC 42001:2023: Information technology—Artificial intelligence—Management system*. <https://www.iso.org/standard/81230.html>
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- National Society of Professional Engineers. (1990). *Use of CADD system* (Board of Ethical Review, Case No. 90-6). [URL ako postoji]
- National Society of Professional Engineers. (2024). *Use of artificial intelligence in engineering practice* (Board of Ethical Review, Case No. 24-2). <https://www.nspe.org/career-growth/ethics/board-ethical-review-cases/use-artificial-intelligence-engineering-practice>
- Organisation for Economic Co-operation and Development. (2024). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- Parker, D. B., Swope, S., & Baker, B. N. (1990). *Ethical conflicts in information and computer science, technology, and business*. QED Information Sciences.
- Romeo, G., & Conti, D. (2026). Exploring automation bias in human–AI collaboration: A review and implications for explainable AI. *AI & Society*, 41, 259–278. <https://doi.org/10.1007/s00146-025-02422-7>
- Torrecilla-Pinero, J. A. (2026). Judgment cannot be delegated: A subject-preserving framework for AI governance. *Philosophy & Technology*, 39, Article 135. <https://doi.org/10.1007/s13347-026-01140-2>
- UNESCO. (2022). *Recommendation on the ethics of artificial intelligence*. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>